

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including gathering and maintaining the data needed, and completing and reviewing the collection of information. Send collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0707-01), Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YY) 28-02-01		2. REPORT TYPE Final		3. DATES COVERED (From - To) 1 April 1998 - 30 Nov 2000	
4. TITLE AND SUBTITLE Enhancing Problem Solving Through Subgoal Learning: The Use of Examples and Computerized Training Environments				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER F49620-98-1-0362	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHORS Catrambone, Richard				5d. PROJECT NUMBER G-42-619	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Georgia Institute of Technology School of Psychology 274 5th Street Atlanta, GA 30332-1070 USA				8. PERFORMING ORGANIZATION REPORT NUMBER Deliverable #1	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research/NL 801 North Randolph St., Rm. 732 Arlington, VA 22203-1977 ATTN: Dr. John F. Tangney				10. SPONSOR / MONITOR'S ACRONYM(S) AFOSR	
				11. SPONSORING / MONITORING AGENCY REPORT NUMBER	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This research has focused on the development of training examples designed to help learners solve novel problems more effectively in domains such as probability and algebra. It has also examined the development of computer-based learning environments, also with the goal of helping learners to acquire knowledge in a way that helps them to transfer to new problems. An additional area of inquiry has been on the development of computer-based tools for carrying out tasks. The common thread among these projects has been the use of a task analysis (or instructional design analysis) to guide the development of the materials and tools. The results of the studies have indicated that materials and tools designed after such an analysis produce significantly better learning and task performance compared to materials that are designed in other ways. Issues concerning how the materials are presented are also important, but seem to be less important than making sure they are based on a task analysis.					
15. SUBJECT TERMS problem solving, instructional technology, human-computer interaction					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 5	19a. NAME OF RESPONSIBLE PERSON Richard Catrambone
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) 404-894-2682

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (AFOSR)
NOTICE OF TRANSMITTAL DTIC. THIS TECHNICAL REPORT
HAS BEEN REVIEWED AND IS APPROVED FOR PUBLIC RELEASE
LAW AFR 190-12. DISTRIBUTION IS UNLIMITED.

FINAL TECHNICAL REPORT

Grant No: F49620-98-1-0362
Period: 1-APR-1998 to 30-NOV-2000
Date of Submission: February 28, 2000
Name of Institution: Georgia Institute of Technology
Title of Project: Enhancing Problem Solving Through Subgoal Learning:
The Use of Examples and Computerized Training Environments
Principal Investigator: Richard Catrambone

Key Findings

- Task (or instructional) analysis provides principled basis for pedagogy, that is, the development of training and testing materials
 - Task analysis reveals subgoals of a problem domain
 - Training examples designed to teach appropriate subgoals improve test performance on novel problems by 10-15% (e.g., Atkinson & Catrambone, 2000)
- Animation in pedagogy improves performance (e.g., Seay & Catrambone, 2001)
 - Learning time not increased with addition of animations
 - Test scores 5-10% higher compared to training supplemented with static images and 10-15% higher compared to training without static images

Overview

There have been four primary thrusts to the research funded by AFOSR. The first has been the development of training examples designed to help learners solve novel problems more effectively in domains such as probability and algebra. The second thrust has been the development of computer-based learning environments, also with the goal of helping learners to acquire knowledge in a way that helps them to transfer to new problems. The third thrust has been on the development of computer-based tools for carrying out tasks. The fourth thrust has been on factors that influence the likelihood of a learner being reminded of a prior problem or story when working on a current one. Papers and presentations based on these efforts are listed in Table 1. Included with this summary are copies of the journal papers and conference proceedings papers listed in Table 1.

The common thread among the first three sets of projects has been on the use of a task analysis (or instructional design analysis; different researchers use different terms; I will use "task analysis" for consistency), to guide the development of the materials and tools. By task analysis I mean an attempt to identify as clearly and objectively as possible the information a person needs to know in order to solve problems or carry out certain tasks in a particular domain. This is a difficult chore and there does not exist a single method for approaching it that is guaranteed to produce the "right" task analysis (my approaches are described below). There are many programs to improve learning but these efforts are often heavy on technology and light on cognitive theory, pedagogy, and experimental methodology. Our work has gone a long way towards improving the design of instructional systems and tools by pushing hard on task analysis and careful experimentation. However, there are two key areas on which additional progress needs to be made.

First, the task analysis approach (e.g., determining what the learner or user needs to know) needs to be more systematized so that a high degree of reliability can be obtained among people doing such analyses. Currently the task analyses require great effort and time by the PI or in some cases his

graduate students and it would be useful to find ways to make the process more streamlined and reproducible so that others are willing to do this kind of analysis.

Second, the research has suggested that while a lot of the variance of instructional quality can be accounted for by the task analysis, the issue of *how* the information is presented also matters. This also includes ways of guiding learners to interact with the information, through computer-based systems, to aid performance. While some of our studies have considered presentation issues, they were usually secondary to the focus on the task analyses and were not systematically examined. Future work will consider systematic examinations of these presentational issues as well as the issue of how to systematize the task analysis.

The rest of this summary focuses on two areas of AFOSR-funded research that I have conducted. These areas were chosen because they illustrate two key issues of my work: 1) the use of task analyses to guide the development of instructional materials, and 2) the examination of how the presentation of information impacts learning.

Helping Learners to Generalize from Examples in Math/Probability: Task Analysis and Instructional Design

It is usually the case that one can not simply (or only) directly tell learners needed information for solving problems in a domain. Students have difficulty grasping such a presentation. Rather, examples, demonstrations, and interactive exercises seem to be essential supplements for helping learners form durable and generalizable solution procedures.

Students prefer to learn from examples, yet they have great difficulty generalizing from examples to novel problems. I make the assumption that learners' problem solving knowledge for a domain can be represented in terms of subgoals and methods for achieving those subgoals. Learners who possess "useful" subgoals are more likely to be able to solve novel problems compared to learners who have acquired only a rote approach to solving problems. Novel problems are often made up of rearrangements of old subgoals and may also require modifications to old methods. Learners who have their knowledge organized by appropriate subgoals and methods have a better chance of making modifications in order to solve novel problems compared to learners who do not have their knowledge organized this way.

For instance, consider the following algebra word problem dealing with work: "Joe can paint a fence in 3 hours. Mary can paint the fence in 2 hours. How long will it take them to paint the fence if they work together?" One can memorize a set of steps for solving this problem and most students have little trouble applying those steps to a new problem that involves Bill and Sue rebuilding a carburetor as long as Bill's rate of work and Sue's rate of work are explicitly stated.

However, complications arise for many students if the new problem does not state Sue's rate directly and mentions that Bill started on the job before Sue (e.g., "Bill can rebuild a carburetor in 4 hours. Sue is twice as fast as Bill. How long will it take them to rebuild a carburetor if they work together and Bill started 1/2 hour before Sue"?). The memorized set of steps does not work in this new case. However, if a student had appropriate subgoals rather than a rote set of steps, he or she might be able to solve the carburetor problem even though he or she had studied only the fence painting example. That is, if the student had learned that in order to solve work problems, one must achieve the "subgoals" of forming a representation for the work rate of each worker and the amount of time each worker works on the task, he or she might be in a better position to figure out how to come up with the right representations to place into the work equation.

My research has supported this conjecture: learners are much more successful solving novel problems in a domain if they have studied examples that were designed to convey the subgoals that

are used to solve most problems in the domain. Even when new problems require new or modified methods for achieving the subgoals, performance tends to be stronger when learners have learned the right subgoals from examples. The subgoals provide a type of scaffolding that guide the learner to the steps that need to be changed. Performance of learners who have acquired subgoals is typically at least 15% better than other learners and in studies currently in progress the differences are even greater.

My typical approach to task analysis in this area (the area of developing examples to help students learn to solve novel problems in mathematics and the "hard" sciences) is to first identify problems I want learners to be able to solve. Next, I solve the problems--either on my own or with the help of a domain expert--and try to identify in as much detail as seems relevant the subgoals I tried to achieve and the steps I used to achieve the subgoals.

Decisions about how "low level" to get in the task analysis itself are determined by a variety of factors including the assumed background of the learners, the time allotted to cover the material, and the aims of the instructor. There will be better and worse materials that one might construct using the task analysis as a guide, but I assert that if the task analysis is *not* done first, then any materials created are less likely to be successful at helping people learn and generalize.

Computer-Based Learning Environments: Algorithm Animation

This line of research examines the use of computer visualization and animation technologies to help teach computer algorithms. Educators who teach computer algorithm design and analysis know what a challenging task it is. The conceptual nature of the domain makes it difficult for students to learn new computer algorithms, understand the algorithms' methodologies, and fundamentally grasp the importance and pervasiveness of algorithmic design throughout the field of computing.

An algorithm animation is a symbolic, dynamic visualization of the data, structures, and operations of an algorithm. The animation can be thought of as a movie of the execution and operation of an algorithm or program. Researchers have hypothesized that this tangible visual presentation will make the abstract algorithmic concepts more concrete, and therefore more understandable to students.

Prior research has been equivocal on this issue. The inconsistent research findings might be due to instructors and researchers not considering carefully enough what it is they want students to learn and how to go about testing it. I conducted a task analysis that has provided, in some sense, the raw material to guide the construction of new animations, the exercises to be done with the animations, and the tests to measure learners' knowledge of the algorithms.

It is instructive to consider how we conducted the task analysis for a particular algorithm, the binomial heap. I, a relative novice, conducted the task analysis by solving a set of "typical" problems dealing with binomial heaps created by my colleague, the domain expert. By asking questions and reading a text, I induced a set of definitions, rules, etc. that enabled me to solve the problems. This was an iterative process. That is, I would use my current state of knowledge to solve a problem and would usually run into an impasse or ambiguity. I would then consult my colleague in order to determine the information needed to get past the impasse. Gradually I developed a set of notes that allowed me to solve the target problems. These notes were the product of the task analysis. It is important to recognize that the information was *not* derived--and probably could not be derived--from a formal analysis of the domain.

We have been examining various factors that influence the impact of animations. These include: 1) the manipulation of animation features that may aid learning the specific bits of information produced by the task analysis; 2) an examination of the effects on learning due to how students

interact with the animations; 3) testing whether individual differences affect how students learn with different styles of animations and different types of interactions with those animations; and 4) a comparison of learning from animations versus learning from static images when the development of both the animations and static images have been guided by a task analysis.

One interesting result from this animation research has been that when learners are required to predict what will happen next in an algorithm--either by indicating what will happen next in the animation or what the next static image will look like--they perform better on test problems compared to learners who did not make predictions. The fact that "prediction" learners using static images did about as well as those who used animations suggests that the prediction element might be what is crucial for aiding comprehension in this domain. However, in other studies we have also found that the animation users completed the learning phase more rapidly and, in some cases, did better on test problems compared to the non-animation users. We have found an approximately 15% performance benefit (accuracy on test problems) by learners who saw animations.

We are beginning to investigate the effects of the learner being able to choose the inputs for the animation rather than having them chosen by the instructor. The ability to choose inputs might help the learner develop his or her understanding of the algorithm more efficiently and more deeply.

Conclusions

The work outlined above and described in detail in the papers accompanying this summary indicate that there is no substitute for a careful analysis of what a learner needs to know or what a user needs to do to carry out tasks. Such an analysis can then guide the development of learning materials and tools. Issues of presentation are also important but can not be productively investigated without a prior task analysis.

Project Publications and Reports

Journal Publications

Stasko, J., Catrambone, R., Guzdial, M., & McDonald, K. (2000). An evaluation of space-filling information visualizations for depicting hierarchical structures. *International Journal of Human-Computer Studies*, 53(5), 663-694.

Byrne, M.D., Catrambone, R., & Stasko, J.T. (1999). Examining the effects of animation and prediction in student learning of computer algorithms. *Computers & Education*, 33, 253-278.

Ram, A., Catrambone, R., Guzdial, M.J., Kehoe, C.M., McCrickard, D.S., & Stasko, J. (1999). PML: Representing procedural domains for multimedia presentations. *IEEE Multimedia*, 6(2), 40-52.

Manuscripts Under Journal Review

Catrambone, R. (under review). *The effects of surface and structural feature matches on the access of story analogs.*

Seay, A.F., & Catrambone, R. (under review). *Using animations to help students learn computer algorithms: A task analysis approach.*

Sukel, K.E., Catrambone, R., Essa, I., & Brostow, G. (under review). *Presenting movement in a computer-based dance tutor.*

Papers in Refereed Conference Proceedings

McCrickard, D. S., Catrambone, R., & Stasko, J. (2001). Evaluating animation in the periphery as a mechanism for maintaining awareness. To appear in *Proceedings of the IFIP TC 13 Conference on Human-Computer Interaction (INTERACT 2001)*. Amsterdam: IOS Press.

Seay, F., & Catrambone, R., (2001). Using animations to help students learn computer algorithms: A task analysis approach. To appear in *Proceedings of the 10th International Conference on Artificial Intelligence in Education*.

Atkinson, R.K., & Catrambone, R. (2000). Subgoal learning and the effect of conceptual vs. computational equations on transfer. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum, 591-596.

McCrickard, D. S., & Catrambone, R. (1999). Beyond the scrollbar: An evolution and evaluation of alternative list navigation techniques. In *Proceedings of the IEEE Symposium on Visual Languages*. Los Alamitos, CA: IEEE Computer Society, 270-277.

Rodenstein, R., Abowd, G., & Catrambone, R. (1999). OwnTime: A system for timespace management. In *Proceedings of CHI Human Factors in Computing Systems Conference (Extended Abstracts)*. New York, NY: Association for Computing Machinery, 200-201.

Catrambone, R. (1998). Generalization by studying examples versus generalization by applying examples to problems. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum, 214-219.

Paper Presentations (if not already listed above as a conference proceeding)

Catrambone, R. (2000). *Studying examples versus solving training problems: It does not seem to matter as long as subgoals are learned*. Paper presented at the 41st Annual Meeting of the Psychonomic Society (New Orleans, LA, November).

Catrambone, R. (1999). *Subgoal learning as a function of formal versus informal labels in examples*. Paper presented at the 40th Annual Meeting of the Psychonomic Society (Los Angeles, CA, November).

Catrambone, R. (1998). *Analogical access: Entities, first-order relations, and higher-order relations*. Paper presented at the 39th Annual Meeting of the Psychonomic Society (Dallas, Texas, November).

Invited Colloquia

Helping Students to Solve Problems in Math and Science: A Task Analysis Approach to Designing Teaching Materials.
Mississippi State, October, 2000.

Creating Better Teaching Materials to Aid Problem Solving in Math and Science.
Grinnell College, Iowa, September, 2000.

Learning Subgoals from Examples to Solve Problems.
University of Saarlandes, Saarbruecken, Germany, August, 2000.

Creating Better Teaching and Training Materials.
The Pennsylvania State University, April, 1999.